

Liquid Cooling v. Air Cooling Evaluation in the Maui High Performance Computing Center

Prepared for the U.S. Department of Energy
Federal Energy Management Program

By Lawrence Berkeley National Laboratory
Rod Mahdavi, PE, LEED AP

July 2014

Contacts

Rod Mahdavi, P.E., LEED AP
Lawrence Berkeley National Laboratory
1 Cyclotron Road
Berkeley, CA 94270
(510) 495-2259
rmahdavi@lbl.gov

For more information on the Federal Energy Management Program, please contact:

Will Lintner, P.E., CEM
Federal Energy Management Program
U.S. Department of Energy
1000 Independence Ave SW
Washington, DC 20585
202-586-3120
william.lintner@ee.doe.gov

Acknowledgements

The author would like to acknowledge the effort of the Maui High Performance Computing Center's David Morton and Joseph Dratz in facilitating this case study.

Abbreviations and Acronyms

°F	Degrees Fahrenheit
CRAH	Computer Room Air Handler
DoD	Department of Defense
DSRC	DoD Supercomputing Research Center
FEMP	Federal Energy Management Program
IT	Information Technology
kW	Kilowatt
MHPCC	Maui High Performance Computing Center
PDU	Power Distribution Units
PUE	Power Usage Effectiveness
SAT	Supply Air Temperature
UPS	Uninterruptible Power Supply

Contents

Executive Summary v

Introduction 6

Direct Water Cooling Option 7

MHPCC Water Cooled System Layout 8

Observations: 9

Conclusions..... 12

List of Figures

Figure 1 – PUE of Water Cooled Riptide HPC vs Air Cooled MANA HPC Systems 7

Figure 2 – Liquid Cooling, IBM iDataplex Server 8

Figure 3 – Riptide’s Direct Water Cooled Chips 9

Figure 4 – Riptide PUE vs IT Load..... 10

Figure 5 – Thermal Maps 11

Figure 6 - Riptide Racks Intake and Exhaust Temperatures During Riptide Linpack Test 12

Executive Summary

This paper describes the efficiency characteristics of a water cooled information technology (IT) system applied in a retrofit project at the Maui High Performance Computing Center (MHPCC) data center. The installation of the water cooled system, Riptide, is the largest computer installation in the history of the MHPCC Department of Defense (DoD) Supercomputing Research Center (DSRC). It places the MHPCC DSRC among the leaders in the DoD research and development community. The MHPCC DSRC leverages world-class technologies to accelerate our nation's ability to meet its most demanding challenges. The new IBM iDataPlex system increases the computational capability of MHPCC's principal HPC platforms to 250+ TeraFLOPS. The study included observation of the data center's physical conditions, the new water cooled computer, the new compressor-less heat rejection system, and environmental conditions. As expected, an evaluation of cooling and electrical system components during system tests showed much less cooling power is required by the water cooled IT system, compared to the cooling power required by the air cooled system. From our testing, we estimate that the water cooling will save \$200,000 per year in operating costs.

Introduction

On average, 50 percent of a typical air-cooled data center's energy consumption and carbon footprint today is not caused by the computing processes themselves, but by powering the necessary cooling systems. When looking at energy efficiency from a holistic perspective, that energy distribution is far from optimal. The heat generated by computing, networking, and storage equipment from the data center was once removed exclusively by moving cooled air through the rack, but increasingly the heat removal is accomplished through liquid cooling. Here, “liquid cooling” involves the application of a direct water cooled system. The use of direct cooling technology represents a paradigm shift in heat removal to one that requires less energy than traditional computer room air conditioning methods. The purpose of our study was to verify the better cooling efficiency of the liquid cooled IT system over the previously used, air cooled IT system. More specifically, we compared the Riptide water cooled system (an IBM iDataplex 360) with MANA (a Dell PowerEdge M610), an air cooled IT system.

Testing was performed just after the Riptide system was installed, but before it took over any of the computing processes from MANA. To run fully loaded IT systems for achieving comparable test results, we performed a Linpack test on both the Riptide and MANA HPC systems. Linpack is a software library for performing numerical linear algebra on digital computers. The Linpack test is designed to assess the performance of the IT system and not its supporting infrastructure. By metering and monitoring power use in the cooling system and losses in the power distribution system during the Linpack test, we could evaluate infrastructure performance fairly across the two systems.

Figure 1 compares the Power Usage Effectiveness (PUE) of the two different systems during tests. The liquid cooling system (Riptide) presented a better PUE (typically 15% better). Measurements showed that 90% of Riptide cooling was through liquid in lieu of air. Cooling for each system was calculated by measuring return air/water temperatures and supply air/water temperature and the amount of air/water flow. In addition, efficiency of cooling, kilowatt per ton (kW/ton), for each system was considered. The IT load was measured at the output of the Power Distribution Units (PDU). Power losses were calculated by deducting the PDU output from the Uninterruptible Power Supply (UPS) input. While water flow was estimated using the pump performance curve, certain estimations were necessary regarding the amount of airflow to each system, namely the number of perforated tiles serving each area. Our measurement showed that the pressure under the raised floor was quite consistent, which made the airflow estimation more trustworthy.

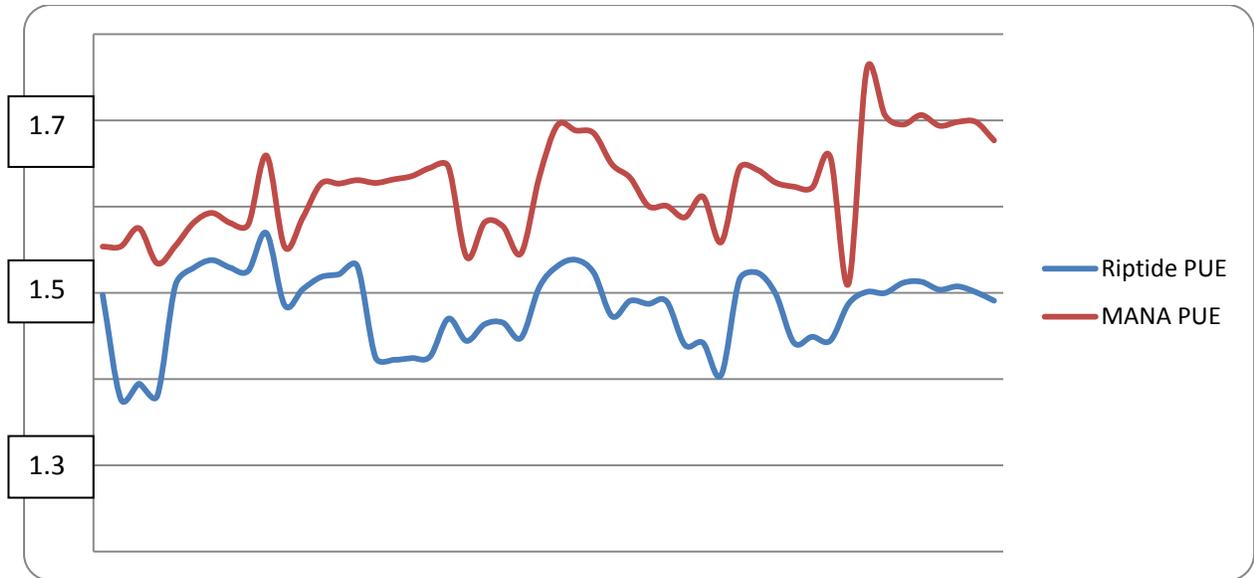


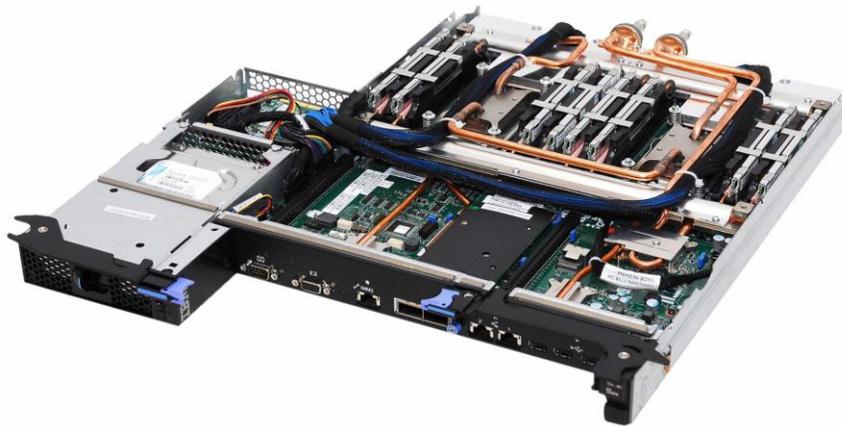
Figure 1 – PUE of Water Cooled Riptide HPC vs Air Cooled MANA HPC Systems

Direct Water Cooling Option

The MHPCC direct water-cooling technology is made possible by the latest IBM iDataPlex server. The processors and other components in the new high performance computer are cooled with inlet water to the server with the cooling water temperature as high as 113 degrees Fahrenheit (°F). In this design, the water cooling loop extends into the servers. Heat is directly dissipated from the microprocessor and memory modules into the cooling water that flows through very fine channels in the heat sink. The hot water coming out of the server is subsequently cooled to a temperature of 113 °F or less as it passes through a dry cooler. This minimizes the need for compressor cooling. In general, using liquid as the working fluid to remove the heat from the server has the following immediate advantages:

- Improved performance per watt
- Lower operating expenses (OPEX) for the data center

Figure 2 illustrates the Riptide computer server with the cooling water piping. Copper pipes carry cooling water to cool plates on the top of the chips.



In 2011, the ASHRAE committee 9.9 published *Thermal Guidelines for Liquid Cooled Data Processing Environments*. This whitepaper introduced five environmental classes specifying operating temperatures and required cooling equipment for each class. Using the Direct Water Cooled iDataplex 360 M4 servers allows the MHPCC data center to run in ASHRAE liquid cool class W4.

Figure 2 – Liquid Cooling, IBM iDataplex Server (Courtesy of IBM)

MHPCC Water Cooled System Layout

Figure 3 illustrates the system layout. Two dry coolers discard heat from the water (top left of the figure). Water then flows to the servers (right bottom of the figure). The water treatment system is shown on the lower left of the diagram. As shown in the middle of the diagram, supply water is piped to the chilled water heat exchangers. If the cooling water supplied from dry coolers is not cool enough then backup cooling is provided by the chilled water system.

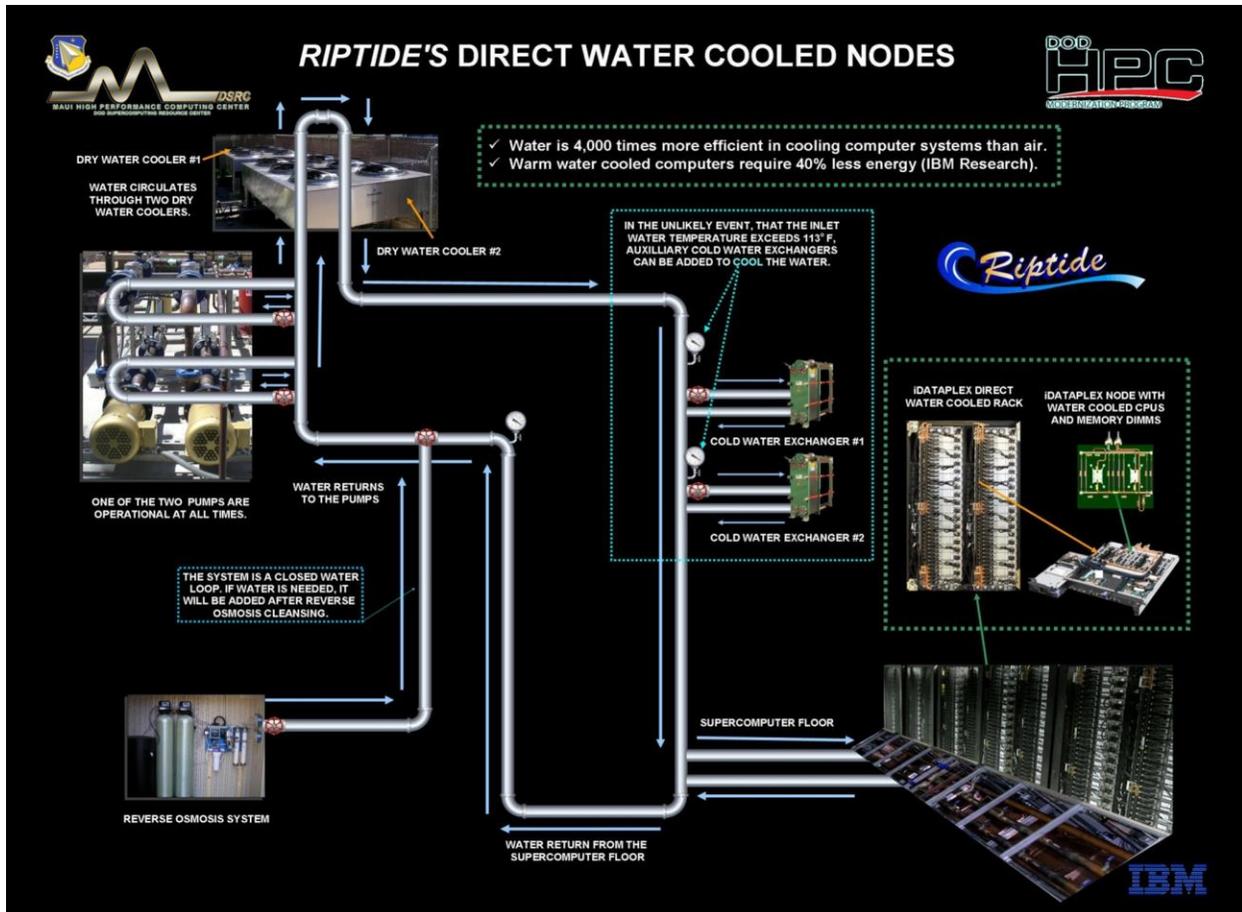


Figure 3 – Riptide’s Direct Water Cooled Chips

Observations:

The following observations were made during the testing period.

- The PUE for the entire data center improved during the Riptide test because of the lower cooling power necessary.
- During Riptide Linpack testing when Riptide was loaded close to 100%, the power required by the dry cooler system was 25kW.
- Riptide dispensed an estimated 15% of the total heat it produced into the air. Considering the efficiency of the chiller plant, that equates to 15kW.
- We concluded that the Riptide water cooling requirement during Linpack testing was about 25kW. This equates to .25kW/ton calculated efficiency for the 70 ton dry cooler, which is much better than the chiller’s efficiency.
- Therefore, 40kW of power is used for cooling the Riptide computer servers.

- Another 10kW was used to cool other Riptide components including storage and network equipment.
- On average, the total cooling power for Riptide was 50kW.
- On average, 130kW was measured to have been used for cooling the MANA system.
- Moreover, as it was expected, lower PUEs were observed when the IT load was higher.
- Figure 4 illustrates the Riptide energy use and the system PUE during the 5 days of the test. The PUE of the system improves substantially from 1.5 to 1.3 when the load on Riptide increases. One reason for the improvement is that variation in cooling is not directly proportional to the IT load. In addition, because of higher load factor, the uninterruptible power supply had a better efficiency with higher IT loads.

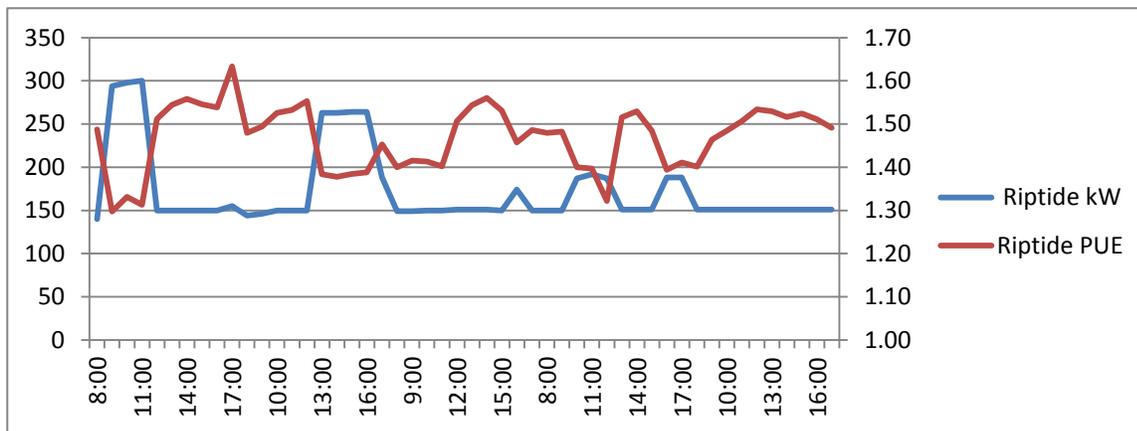


Figure 4 – Riptide PUE vs IT Load

Figure 5 illustrates thermal maps of the data center during different IT loading. Temperature measurement is done at about 6 feet above the floor. Figure 5A shows Riptide loaded and rejecting heat into its hot aisle (10-15% of cooling was done by air). Figure 5B shows Riptide not loaded and most of the heat generated by Riptide is removed by water. Figure 5C shows the MANA hot aisle high temperatures when MANA was loaded to maximum and Riptide had an average load. Figure 5D illustrates the change to the thermal map when MANA was unloaded.

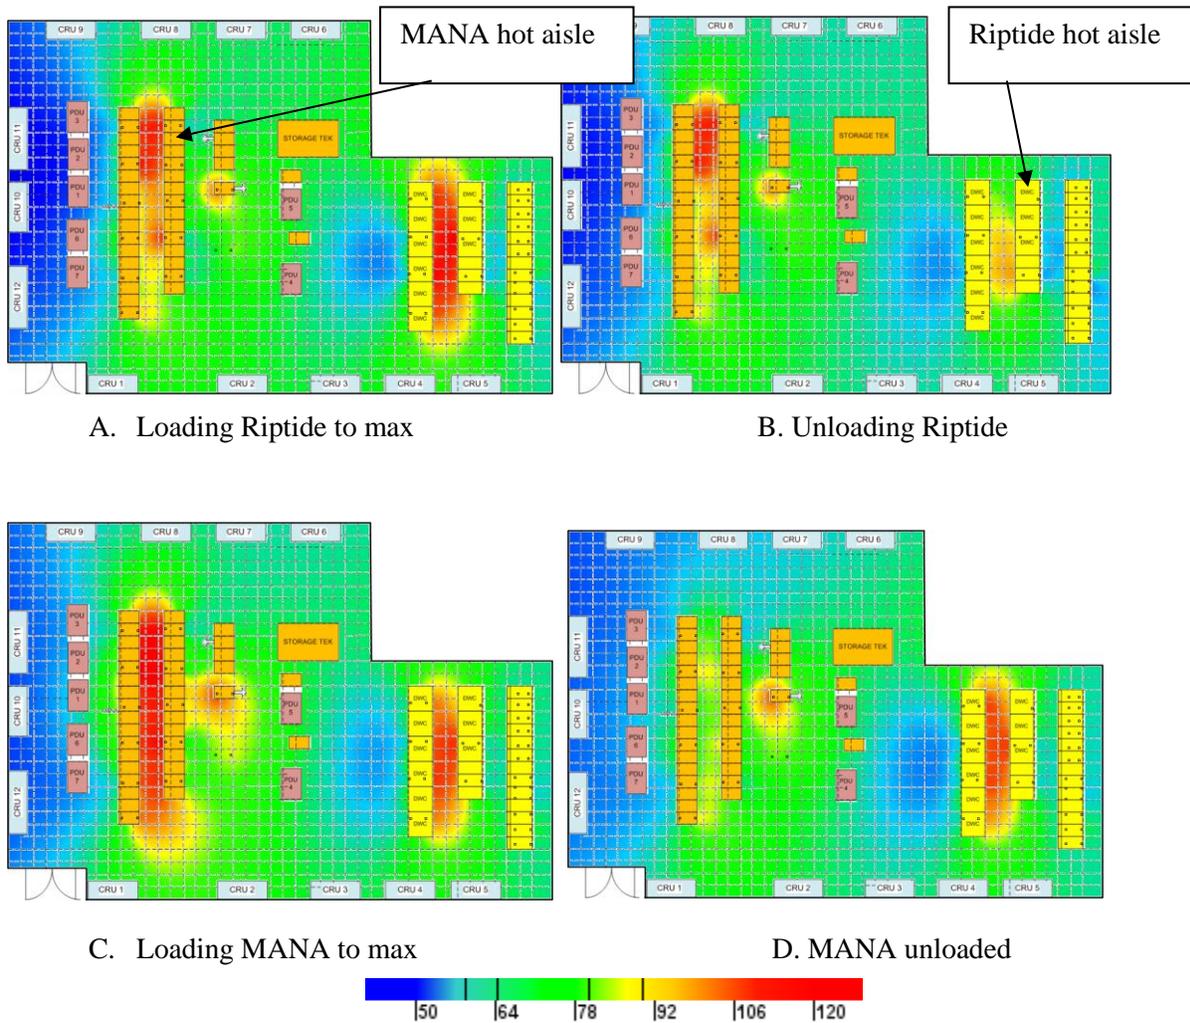


Figure 5 – Thermal Maps

Riptide Air Cooled Load:

During Riptide Linpack test, the Riptide rack exhaust air temperatures increased (about 5⁰F) as illustrated in Figure 6 during the Riptide Linpack test. Notice that rack air intake temperatures were increased also and that is because computer air handlers delivering air for cooling reacted with a delay to the new higher cooling load.

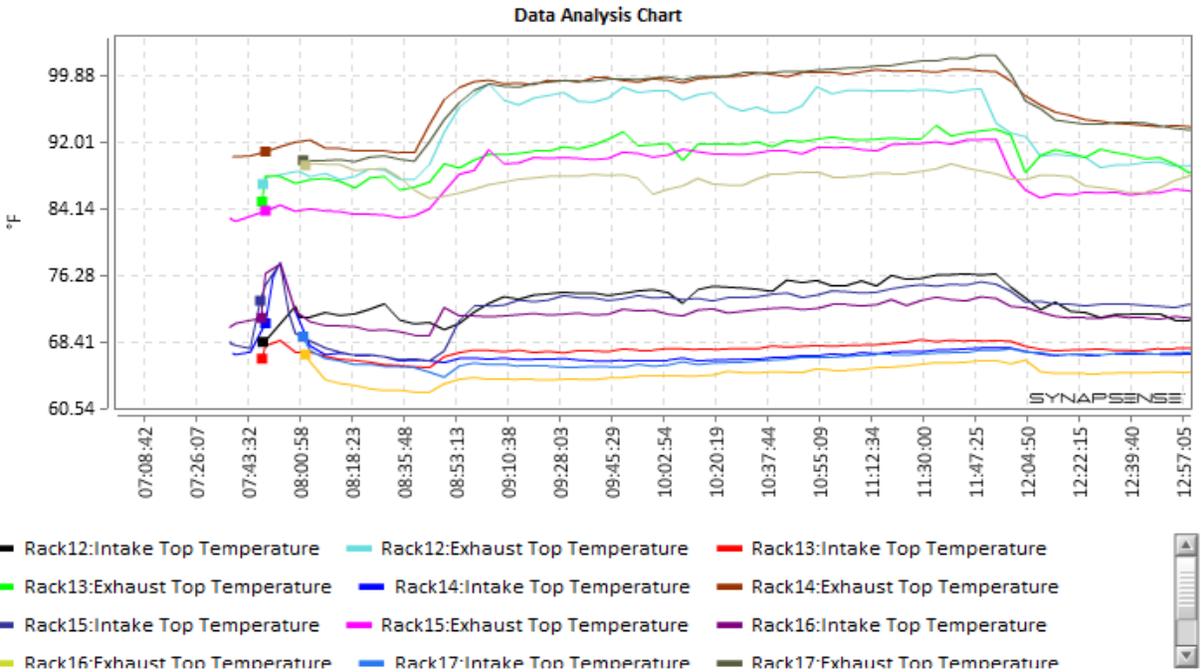


Figure 6 - Riptide Racks Intake and Exhaust Temperatures During Riptide Linpack Test

Conclusions

The heat generated by computing and storage equipment from the data center was once removed exclusively by moving chilled air through the rack, but increasingly that removal is accomplished by liquid cooling. In the case of the Riptide HPC system, liquid cooling involved the application of a direct water-cooled system. The use of direct cooling technology represents a paradigm shift in heat removal to one that requires less energy than traditional computer room air conditioning methods. We estimated a potential savings of \$200,000 in operating costs per year for the 300kW MHPCC DSRC system when using a direct water cooling system instead of an air-cooled system.

U.S. DEPARTMENT OF
ENERGY | Energy Efficiency &
Renewable Energy

For more information, visit
femp.energy.gov

DOE/EE-1119 • July 2014